

<https://blog.wenxuecity.com/myblog/80392/202304/17936.html>

人工智能 AI A Modern Approach 人工智能学：一种现代方法

(2023-04-19 19:08:37)

人工智能学 AI: A Modern Approach 人工智能学：一种现代方法

人工智能学：一种现代方法 第一章 逻辑和知识表示

人工智能学：一种现代方法 第二章 SAE 人工智能学：一种现代方法 第二章 SAE

AlphaGo 人工智能学：一种现代方法 第三章 AlphaGo 人工智能学：一种现代方法 第三章 AlphaGo

人工智能学：一种现代方法 第四章 Logical positivism or logical empiricism 人工智能学：一种现代方法 第四章 Logical positivism or logical empiricism

人工智能学：一种现代方法 第五章 逻辑和知识表示 人工智能学：一种现代方法 第五章 逻辑和知识表示

人工智能学：一种现代方法 第六章 逻辑和知识表示 人工智能学：一种现代方法 第六章 逻辑和知识表示

人工智能学：一种现代方法 第七章 preferences Preferences 人工智能学：一种现代方法 第七章 preferences Preferences

人工智能学：一种现代方法 第八章 preference 人工智能学：一种现代方法 第八章 preference

人工智能学：一种现代方法 第九章 AI: A Modern Approach 人工智能学：一种现代方法 第九章 Turing Test 人工智能学：一种现代方法 windtunnel approach 人工智能学：一种现代方法 windtunnel approach

人工智能学：一种现代方法 第十章 Artificial general intelligence 人工智能学：一种现代方法 Artificial general intelligence

人工智能学：一种现代方法 第十一章 Artificial general intelligence 人工智能学：一种现代方法 Artificial general intelligence

人工智能学 AI: A Modern Approach 人工智能学：一种现代方法 第十二章 人工智能学

---

Page 3: Turing deserves credit for designing a test that remains relevant 60 years later.

Page 5: The standard of rationality is mathematically well defined and completely general, and can be “unpacked” to generate agent designs that provably achieve it.

Page 5: Aristotle (384–322 B.C.), whose bust appears on the front cover of this book, was the first to formulate a precise set of laws governing the rational part of the mind.

Page 25: In terms of methodology, AI has finally come firmly under the scientific method. To be accepted, hypotheses must be subjected to rigorous empirical experiments, and the results must be analyzed statistically for their importance (Cohen, 1995).

Page 483: Preferences, as expressed by utilities, are combined with probabilities in the general DECISION THEORY theory of rational decisions called decision theory:

Decision theory = probability theory + utility theory .

The fundamental idea of decision theory is that an agent is rational if and only if it chooses the action that yields the highest expected utility, averaged over all the possible outcomes MAXIMUM EXPECTED of the action. This is called the principle of maximum expected utility (MEU). Note that Utility "expected" might seem like a vague, hypothetical term, but as it is used here it has a precise meaning: it means the "average," or "statistical mean" of the outcomes, weighted by the probability of the outcome.

Page 612-613: To address this issue we list six constraints that we require any reasonable preference relation to obey: Orderability..., Transitivity..., Continuity..., Substitutability..., Monotonicity..., Decomposability...

These constraints are known as the axioms of utility theory. Each axiom can be motivated by showing that an agent that violates it will exhibit patently irrational behavior in some situations.